SEEGERWEISS LLP
NEW YORK • NEW JERSEY • PHILADELPHIA

LEVI&KORSINSKY LLP
New York / Washington D.C. / Connecticut / California

November 20, 2018

VIA ECF
Hon. Michael H. Simon, U.S.D.J.
United States District Court
District of Oregon
Mark O. Hatfield U.S. Courthouse, Room 1527
1000 Southwest Third Avenue
Portland, Oregon 97204

In re: Intel Corp. CPU Marketing, Sales Practices and Products Liability Litigation, Case No. 3:18-md-02828-SI

Dear Judge Simon:

On behalf of the Plaintiffs' Steering Committee, we respectfully submit this letter in connection with the parties' Joint Discovery Report. As discussed in the Joint Discovery Report and detailed below, there are several discreet issues that remain in dispute.

At present, no documents have been produced by Intel to Plaintiffs. As the Court expressed at the initial June 15, 2018 conference, it had contemplated that Intel would produce relevant documents prior to resolution of Intel's Rule 12(b)(6) motion. Despite nearly 6 months since the parties' initial conference, we are informed that Intel has once again not even begun its review of relevant documents for production, and now after months of delay seeks a stay of discovery again.

I. THE COURT SHOULD IMPOSE A DEADLINE FOR THE EXCHANGE OF RULE 26(A)(1) DISCLOSURES

As contemplated by Fed. Rule of Civ. Proc. 26(a)(1) and Local Rule 26-1, Plaintiffs have requested the identity of Intel's custodians who possess relevant information, custodial and non-custodial data sources, and the nature and scope of electronically stored information ("ESI") since at least June 2018 when the parties began conferring about appropriate pretrial orders. *Inter alia*, Rule 26(a)(1) provides the parties shall disclose:

- The name and, if known, the address and telephone number of each individual likely to have discoverable information—along with the subjects of that information—that the disclosing party may use to support its claims or defenses; and

- A description by category and location of all documents, electronically stored information, and tangible things that the disclosing party has in its possession, custody, or control and may use to support its claims or defenses.

This information is critical to timely and efficiently conduct discovery and thus generally is required to be automatically disclosed at the outset of the case.

On August 10, 2018, Intel disclosed that it had sent litigation hold notices to over 500 custodians believed to have knowledge relating to Plaintiffs' allegations, and, further, that it has worked with these custodians to identify key data sources beyond individual custodial files. The legal hold notices were sent to custodians as early as January 2018, nearly one-year ago now. In particular, Intel advised that the custodians fell within the following specific categories of relevant information:

- Discovery of the side channel attacks and development of mitigations;
- Microarchitecture processor design dating back to 1995;
- Security vulnerabilities; and
- Advertising, marketing, pricing and sales information and data.

Annexed hereto as Exhibit 1 is a copy of an email chain between the parties regarding preservation and disclosure of custodians, data sources, and relevant categories of discoverable information and data.

Significantly, Intel identified these custodians, categories of relevant information, and data sources prior to Plaintiffs filing their Consolidated Class Action Complaint ("Complaint") [Dkt. No. 115] and prior to receiving Plaintiffs' First Request for Production of Documents ("RFPs"). After numerous requests, Intel has now agreed to disclose the list of custodians who received litigation holds and provided information about the relevant knowledge they possess by December 7, 2018. Intel still has not disclosed the custodial and non-custodial data sources identified as containing relevant information – despite Intel's agreement to do so.

Pursuant to Pretrial Order No. 4 (Order Establishing Protocol for Document Collection and Production), the parties were required to meet and confer on (a) the identity and role of custodians possessing relevant information from whom documents will be collected and produced; (b) search methodology and search terms, if any, to be applied, and the use of technology assisted review ("TAR") or similar technologies; (c) any locations and descriptions of relevant data sources, including custodial, non-custodial, and third-party documents; and (d) any applicable and appropriate timeframes for the collection, review, and production of documents.

On November 15, 2018, 11 months after Intel sent its legal hold notices and four months after Intel first disclosed to Plaintiffs that it had identified over 500 custodians with relevant knowledge, Intel served Plaintiffs with a list of 38 custodians it intends to use for purposes of composing the corpus of documents to review. Intel designated the list confidential pursuant to Pretrial Order No. 5. Intel has not provided sufficient information to discern what relevant information these Group I custodians possess and over what period of time – especially given the allegations in the Complaint go back to 1995.

During the parties' meet and confer conferences, Intel confirmed that there were many additional custodians who possess relevant knowledge and from whom Intel has collected documents and/or ran key word searches. Despite months of conferring with Intel, as Plaintiffs' efforts to informally obtain this information have been protracted, it is necessary for the Court to require Rule 26(a)(1) disclosures in order to facilitate a timely and efficient discovery process. Without disclosures, Plaintiffs will not have necessary information to provide any meaningful feedback or approval of the custodians from which Intel has collected documents.

To the extent necessary in pursuit of the details, Plaintiffs are prepared to pursue a corporate structure and ESI 30(b)(6) deposition to identify relevant employee custodians, teams and departments, and learn about the manner and methods used by Intel to store and maintain electronically stored information.

II. TECHNICAL INFORMATION CONCERNING THE ALLEGED DEFECT, INCLUDING UNDERLYING DESIGN, DEVELOPMENT, AND ENGINEERING DOCUMENTS REGARDING INTEL'S IMPLEMENTATIONS OF SPECULATIVE EXECUTION, MEMORY ACCESS PROTECTION, AND SUPPORTING TECHNOLOGIES IN ITS CPUs IS RELEVANT TO PLAINTIFFS' CLAIMS AND SHOULD BE PRODUCED

In their First Request for Production of Documents ("RFPs") served on October 3, 2018, Plaintiffs seek production of technical information concerning the alleged defect, including underlying design, development, and engineering documents regarding Intel's implementations of speculative execution, memory access protection, and supporting technologies in the subject processors, which Plaintiffs allege gives rise to the vulnerability. *See* Exhibit A to Joint Discovery Report, Plaintiffs' RFP Nos. 2-4, 9, 11. These documents are central to all of Plaintiffs' claims and go to the heart of Intel's liability. Intel, however, objects to Plaintiffs' discovery in a seeming attempt to limit responsive information to simply its knowledge of the exploits known as Spectre, Meltdown and Foreshadow. Despite Intel's objection, however, the requested discovery is relevant to Plaintiffs' allegations and should be produced.

A. Intel's unilateral limitation on the scope of discovery.

Intel objects, on grounds of relevancy, to the production of technical documents concerning the design, development, and engineering of Intel's implementations of speculative execution and memory access protection. In large part, Intel does this by re-defining and limiting the defect at issue. More specifically, in the RFPs, Plaintiffs define the "defect" to mean the security vulnerability created when Intel's CPUs engage in speculative execution, causing information that should remain secure from being rendered accessible to unauthorized use, including, but not limited to,

a. Implementing a shared Cache design that does not (i) include any mechanism to ensure that sensitive or privileged information (or data concerning that information) was flushed once the processor determined it had unnecessarily or improperly

accessed the Cache or (ii) provide any protection against Side Channel Attacks that access data that remaining in the Cache;

b. Implementing Speculative Execution in its CPUs such that it created a window of time during which an attacker could make unnecessary and unauthorized requests to access the Cache for information; and

c. The failure of its CPUs to flush unauthorized (e.g., exceptions) and unnecessary (e.g., mistakes) memory requests until such time as the processor was ready to retire the instructions in program order.

Plaintiffs' definition of the defect aligns with Plaintiffs' allegations in the Complaint. *See* Complaint, ¶¶ 2, 180, 222-223, 260.

In its Responses, Intel objects to Plaintiffs' definition of "defect" because Intel asserts its "CPUs are not defective" and only the Meltdown, Foreshadow, and Spectre exploits "are at issue in this litigation." Intel's Responses are, thus, "limited to those security vulnerabilities." *See* Exhibit B to Joint Discovery Report, Intel's Response Nos. 2-4, 9, 11. Plainly, Intel's narrow interpretation of the "defect" as alleged by Plaintiffs to mean something entirely different than what is alleged is untenable. That is because the defect as described in Plaintiffs' Complaint and RFPs expressly encompasses Intel's underlying design, development, and engineering decisions in connection with Intel's implementations of speculative execution, memory access protection, and supporting technologies in Intel's CPUs. Despite Intel's characterization, the security vulnerability or defect in Intel's CPUs is not Meltdown, Foreshadow, or Spectre. Rather, those side channel attacks exploit the security vulnerability that existed in Intel's CPUs to access privileged data.

Intel also argues that Plaintiffs' discovery is irrelevant because Plaintiffs' claims "sound[] in fraud." *Id.* Intel is wrong. The security vulnerability in Intel's CPUs arises out of Intel's design and implementations of speculative execution and memory access protection, and these technical documents go to the core of Plaintiffs' allegations. It was these undisclosed Intel choices that allowed the exploits to succeed.

The Complaint makes plain that Intel's knowledge and decisions concerning the design and development of the subject chips are at issue. In the Complaint, Plaintiffs allege that the attacks identified in 2018, dubbed Meltdown, Spectre, and Foreshadow, exploited undisclosed vulnerabilities in Intel's CPU design. Complaint ¶¶ 2, 167, 216. Plaintiffs allege that Intel concealed material information regarding the defects when selling its CPUs that Intel knew to be defective. *Id.* ¶ 346. Specifically, Plaintiffs claim Intel knew of its flawed implementations of speculative execution and memory access protection that made information, which should have remained secure and inaccessible to unauthorized use, accessible to unauthorized use. *Id.* ¶¶ 2, 6, 7, 222, 244-246, 248. Intel's overly aggressive implementations created a vast security vulnerability that could be accessed through a number of different exploits including Meltdown, Foreshadow, and Spectre. *Id.* ¶¶ 2, 255-282. Plaintiffs claim that Intel knew, or was reckless in not knowing, about these vulnerabilities. *Id.* ¶ 233, 254.

For example, the Meltdown exploit got its name due to its ability to effectively dissolve Intel's informational barrier that protects privileged data, allowing an attacker to read sensitive information like passwords, login keys, and encryption keys. *Id.* ¶ 255. Meltdown takes advantage of the fact that Intel CPUs use speculative execution to fetch data before enforcing a privilege check to confirm that the user is authorized to read such data. *Id.* ¶ 256. Plaintiffs allege Intel purposefully sacrificed its CPUs' security and implemented speculative execution and memory access protection this way to make Intel CPUs faster. *Id.* ¶ 3, 222. Even though all major CPU developers and manufacturers implement speculative execution, only Intel CPUs suffer from the Meltdown exploit because only Intel implemented speculative execution and memory access protection this way. *Id.* ¶ 265. Like Meltdown, the Foreshadow exploits are based on the fact that Intel's hardware speculates past permission checks, allowing a malicious process a window of time during which it can steal sensitive information. *Id.* ¶ 266.

It is Plaintiffs' position that Intel's failure to maintain the security of privileged information during speculative execution created the security vulnerability that Meltdown, Foreshadow, and Spectre exploited. *Id.* ¶¶ 219, 223, 245, 255-257, 266, 272, 276, 280, 281. Stated differently, the Meltdown, Foreshadow, and Spectre exploits are impossible without Intel's intentional and overaggressive implementation of speculative execution and associated memory protection. For example, as explained in paragraph 256 of the Complaint, Intel CPUs use speculative execution to fetch data before enforcing a privilege check to confirm that the requesting user is authorized to see such data. Similarly, as explained in paragraph 270 of the Complaint, Intel CPUs were designed so that if the address translation process is prematurely terminated through a page fault, the L1 cache lookup is still performed based on the physical address pointed to in the page table (which is no longer the physical memory of the requesting process), and speculative instructions are temporarily permitted to perform computations using privileged data that the process is unauthorized to access from the cache. In both these instances, because of the undisclosed defects, Intel CPUs improperly leave privileged information unprotected during speculative execution. *Id.* ¶ 245.

Intel's overaggressive implementations of speculative execution, memory access protection, and supporting technologies trace back to when Intel first implemented speculative execution in a manner that exposes privileged information to non-privileged access. *Id.* ¶ 5. While the *disclosure* of these defects occurred in or about 2017/2018 in connection with the discovery of Meltdown, Foreshadow, and Spectre, these vulnerabilities themselves were incorporated into Intel's CPU design far earlier – as early as 1995 – when Intel first introduced speculative execution into its CPU design and associated memory protection. *Id.* ¶ 263. Accordingly, the design, development, and engineering of speculative execution and associated memory protections since 1995 is highly relevant and may lead to the discovery of admissible evidence.[1]

[1] Courts that have considered the discoverability of design, development, and engineering documents in the context of allegations of a defendant's fraudulent omission of a known defect hold that Plaintiffs are entitled such technical documents in discovery. *See, e.g., Catalano v. BMW of N. Am., LLC*, No. 15-CV-4889 (KBF), 2016 WL 3406125, *7 (S.D.N.Y. June 16, 2016) (rejecting the defendant's arguments, the court found that plaintiffs' requests for documents and information relating to design and testing are highly material and "vital to allow [plaintiff] to prove

Inter alia, Plaintiffs requested technical documents – together with documents reflecting communications and discussions thereof – related to the development, design, or engineering of Intel's implementation of speculative execution in Intel's CPUs, including:

- Copies of the microarchitecture specifications (MAS) and each redbook related to Intel's defective CPUs;

- Each version of register-transfer language (RTL) code or hardware description language (HDL) code, such as Verilog HDL, VHSIC HDL, or Intel internal HDL, used to model, synthesize or construct any defective CPU; and

- Testing, validation, or other study involving speculative execution in the defective CPUs or the defect.

Plaintiffs need these documents to prove their claims at trial. Much of what Plaintiffs know about the likely cause of the security vulnerabilities at issue here are necessarily based on opinions of experts and what has been written about Intel's CPU. The foundation of that information is limited to what Intel has elected to publicly disclose. To know the precise algorithm that Intel operationalized in the chips with regard to speculative execution and memory protection require production of Intel's microarchitecture specifications (MAS), RTL source code, microcode, the revision history of these sources, and other technical documents including how security issues are handled. The documents will allow Plaintiffs to determine what was done and when with regard to the vulnerability. The source code, in particular, provides irrefutable evidence of how the hardware in Intel CPUs operates. Documents – especially internal communications, memos, and design history changes – related to the design and development of speculative execution will also inform Plaintiffs on whether the security impact of deferred access privilege checks was ever considered, and the extent to which Intel engineers were aware of the potential pitfalls.

By limiting discovery to Meltdown, Spectre and Foreshadow, Intel necessarily obstructs Plaintiffs' discovery of Intel's implementations of speculative execution, memory access protection, and supporting technologies in its CPUs and correspondingly Intel's knowledge of the security vulnerability created thereby. There are many ways to implement speculative execution and not all are susceptible to cache side channel attacks.

---

that the defects in fact existed so that he can establish that [d]efendants' omissions were fraudulent); *Cholakyan v. Mercedes-Benz USA, LLC*, No. CV 10-5944 MMM(JC), 2012 WL 12878362, *9 (C.D. Cal. Apr. 3, 2012) (consumers brought fraudulent omission claims against a car manufacturer, alleging that the defendant knew at the time of sale that its vehicles contained design and/or manufacturing defects that caused water leaks and flooding but concealed it; reaffirming its earlier order, the court required the defendant to all documents relating to the "original design," and all "proposed or implemented modifications thereof," for the water drainage system).

B. The requested documents to which Intel objects are "proportionate" to the needs of this litigation.

Intel also objects on the basis of proportionality to producing technical information concerning the alleged defect, including underlying design, development, and engineering documents regarding Intel's implementations of speculative execution, memory access protection, and supporting technologies in the subject processors, which Plaintiffs allege gives rise to the vulnerability. Intel's Responses Nos. 2-4, 9, 11.

Intel, however, cannot refuse discovery simply by making a boiler plate objection that the discovery is not proportional to the needs of the case; it must explain why, using the factors identified in Rule 26 of the Federal Rules of Civil Procedure. Rule 26 defines the scope of discovery in federal district courts, explaining that, "[p]arties may obtain discovery regarding any non-privileged matter that is relevant to any party's claim or defense and proportional to the needs of the case." Fed. R. Civ. P. 26(b)(1). Rule 26 identifies the following six factors the Court should consider in making a proportionality determination: 1) "the importance of the issues at stake in the action," 2) "the amount in controversy," 3) "the parties' relative access to relevant information," 4) "the parties' resources," 5) "the importance of the discovery in resolving the issues," and 6) "whether the burden or expense of the proposed discovery outweighs its likely benefit." *Id.*

The Committee Notes to the 2015 Amendments make clear that the proportionality Rule does not "permit the opposing party to refuse discovery simply by making a boilerplate objection that it is not proportional." 2015 Committee Notes to Fed. R. Civ. P. 26. Although Rule 26 does not address who bears the burden of establishing proportionality, the Committee Notes make clear that the requesting party does not need to make a threshold showing of proportionality before making a discovery request. *See id.* (explaining that the proportionality requirement "does not place on the party seeking discovery the burden of addressing all proportionality considerations"). The Committee Notes explain that, in addressing the proportionality factors, the burden is on the party that is in the best position to address the particular factor. *See id.* ("A party claiming undue burden or expense ordinarily has far better information -- perhaps the only information -- with respect to that part of the determination."); *see also First Niagara Risk Mgmt., Inc. v. Folino*, 317 F.R.D. 23, 28 (E.D. Pa. 2016) (explaining that "the burden is on [the requesting party] to show that the material it requests is relevant to its claims, and, if that burden is met, [the opposing party] must show that the factors in Rule 26 weigh in favor of our denying [the requesting party's] request for otherwise relevant information").

Here, while Intel does not even attempt to explain how such discovery is not proportional to the needs of the case, all of Rule 26's factors weigh heavily in Plaintiffs' favor. The issues discussed above and the documents sought by Plaintiffs are highly relevant to the claims and defenses at issue. As described above, Plaintiffs allege that Intel knew of the security vulnerability in its CPUs long before the side channel attacks – Meltdown, Foreshadow, and Spectre – were publicly disclosed. The substantial issues relating to Intel's design, development, and engineering decisions in connection with Intel's implementation of speculative execution, memory access protection, and supporting technologies go back as far as 1995. Such information

is highly relevant and internal discussions of these subjects, whether contained in emails, memoranda, or presentations, should be produced, together with the technical documents.

It is also undisputed that the amount in controversy in this MDL, which concerns millions of affected CPUs, is in the billions. Though information and documents are solely in the hands of Intel, Intel has the financial wherewithal to perform the collection, review, and production for the entirety of the relevant time period. Notably, Plaintiffs have agreed to a number of tools – most prominently, deduplication and TAR – to minimize the volume of material to review.

The requested information is relevant and proportionate to the needs of this case. Intel should be so directed, and the subject discovery promptly produced. *See* RFPs 2-4, 9, 11.

III. INTEL SHOULD BE DIRECTED TO COMMENCE ITS TAR REVIEW ON A ROLLING BASIS AND NOT USE SEARCH TERMS BEFOREHAND

There are two primary disputes between the parties concerning the proposed TAR protocol and Intel's review and production. The first concerns the timing of Intel's review and production, while the second involves the use of search terms in conjunction with TAR. Annexed hereto as Exhibits 2 and 3 are copies of Intel's TAR protocol and Plaintiffs' comments to Intel's TAR protocol, respectively.

A. Intel should start training the TAR tool and produce any documents not in dispute.

Intel has been toying with using TAR ever since it filed its supplemental letter to the Court on the parties' pretrial order submissions [Dkt. No. 120]. Despite engaging a consultant, Maura R. Grossman, and performing numerous "road testing" on the TAR tool, Intel has only committed to using TAR for the 38 custodians disclosed to date (and possibly several others undisclosed or unknown).

But even then, it is Intel's position that it will not begin training the TAR tool until after the parties (with the Court's assistance) resolve all of Intel's objection to Plaintiffs' Group I RFPs. As described above, Intel has asserted substantial objections to the scope of discovery. However, there is a set of materials that Intel does not dispute is relevant and responsive. There is no reason not to begin the TAR process and producing those materials.

It is approaching December, and Plaintiffs' opposition to Intel's Rule 12(b)(6) motion is due on December 14, 2018. As expressed at the parties' June 15, 2018 status conference, the Court contemplated that Intel would have made a production by now – in advance of the briefing on the Rule 12(b)(6) motion. Despite many months and many meet and confer conferences with Intel, no documents have been produced by Intel to Plaintiffs.

There is no legitimate reason to delay TAR review and production. Intel can educate the TAR tool using the responsive documents not in dispute, and continue to educate the tool as objections are resolved as to the scope of permissible discovery. This continuous learning will not

impact validation. The only justification asserted by Intel is the burden of training the TAR tool over time. But Intel's argument is untenable. Intel has already sought to serve objections to Group II RFPs after the Court's decision on its Rule 12(b)(6) motion instead of serving now and resolving any objection to the RFPs, retrain the TAR tool, and review documents from custodians that may overlap between Group I and Group II RFPs. . Intel is already contemplating training the TAR tool over time. Having voluntarily taken such an approach to discovery, moving forward with discovery now would not be prejudicial to Intel. Plaintiffs are entitled to production of relevant and responsive documents that are not in dispute. Indeed, such documents should have already been produced.

B. Intel should not be permitted to use search terms before performing a TAR review

Also troubling is that Intel's suggestion to reserve the option to cull documents using key word searches first, before running TAR. Even Intel's consultant, Ms. Grossman, has opposed using search terms prior to TAR as Intel has proposed.

As explained by Plaintiffs' consultant, Doug Forrest, in connection with the parties' submission on the parties' ESI protocol [Dkt. No. 118-2], in contrast to the 80% responsiveness of TAR techniques, key word searches identify only 20-25% of responsive documents. If Intel used key word searches prior to and in conjunction with TAR, then Intel would collect 20-25% of responsive documents. Then Intel would perform TAR on the 20-25% resulting in a production to Plaintiffs of 80% of the 25% or less than 20% of Intel's total responsive document population. That is simply not acceptable. Intel should not be permitted to use search terms in conjunction with TAR.

Use of search terms prior to TAR is an improper methodology to cull and then compose a corpus of documents (proportionally to increase richness of relevant documents) to which TAR will be applied, and therefore precision. As Pretrial Order No. 4 recognizes: "the Parties will meet and confer regarding reasonable and appropriate methods to increase the relative precision or proportion of relevant and responsive documents within the search results and production sets." Pretrial Order No. 4, ¶ 4.

Ms. Grossman and her partner, Gordon V. Cormack, have attacked this improper methodology in an article, "*The Implications of Rule 26(g) on the Use of Technology-Assisted Review*" (FCLR, 7:1, 2014) (annexed hereto as Exhibit 4), taking issue with an earlier article, "*The Implications of Rule 26(g) on the Use of Technology-Assisted Review*" (FCLR, 7:1, 2013). The earlier article suggested that, "if done carefully and correctly, counsel can pre-cull the ESI collection with keywords to improve its richness." Ms. Grossman emphatically disagreed with the authors on the use of search terms prior to TAR:

> As illustrated in this section, an obligation to enrich the collection would increase the complexity and cost of the review, while providing only the illusion of improved validation.… In order to enrich the collection, it would be necessary to discard the vast majority of the documents, at least some of which would almost certainly be responsive and, by virtue of being discarded prior to the TAR effort,

> never reviewed or produced.… *As a consequence, it is entirely possible, as illustrated below, that a review effort achieving a higher recall on the enriched collection might actually find fewer responsive documents overall---and incur a higher level of uncertainty as to the quality of the result---than a review effort achieving a lower recall on the original collection*.

Id. at 293 (emphasis added). In a hypothetical, Ms. Grossman states that TAR by itself has a 70% recall rate while the recall of TAR preceded by keywords is only 56%. Ms. Grossman makes the further point that calculation of recall in both instances has to be done by sampling the same original collection, i.e., the one that existed prior to keywords: "A valid comparison demands that both recall values be derived from the original collection, in which case we see that, in this simple example, the recall of TAR alone is 70%, while the recall of enrichment plus TAR is 56%." *Id.* at 294.[2]

> Ms. Grossman posed the following hypothetical:
>
> Even if each phase, alone, excludes relatively few responsive documents, the combined effect can be substantial. Let us suppose in our hypothetical matter that the enrichment process were to discard 30% of the relevant documents (i.e., were to have a recall of 70%), the TAR process were to have a recall of 75%, and the final human review were to have a recall of 70%. These numbers, in isolation, might be considered reasonable, but consider the combined effect of all three phases on our example. Of 10,000 responsive documents in the original collection, 7,000 (i.e., 70%) would be retained in the enriched collection. Of the 7,000 responsive documents in the enriched collection, 5,250 (i.e., 75%) would be retained in the review set. Of the 5,250 responsive documents in the review set, assuming that none were withheld as privileged, 3,675 (i.e., 70%) would be identified for production. *Accordingly, the recall of the end-to-end review effort would be 36.75% (i.e., 3,675 of the original 10,000 responsive documents*). *Id.*
>
> Simply put, key word searches should not be used prior to or in conjunction with TAR.

---

[2] Relatedly, Ms. Grossman has also opposed TAR validation based upon sampling a key word population, as Intel proposed. Ms. Grossman has advocated for "end-to-end" validation of the entire process beginning once the original collection has been defined and ending after manual review:

> Validation, we argue, should apply to the end-to-end review, starting with the original collection and ending with the production set, regardless of which, if any, of the steps are deemed to be "TAR." *That is, validation must account for all responsive documents excluded by the review, whether before, during, or after "TAR"; or even when traditional review methods are applied*.

Id. at 300 (emphasis added).

C. INTEL'S RENEWED EFFORT TO STAY DISCOVERY SHOULD BE DENIED

Intel had already requested a stay of discovery in its initial case management proposal [Dkt. 83]. At the June 15, 2018 conference, the Court rejected Intel's request and advised the parties that it would not stay discovery pending Intel's motion to dismiss. Despite the Court's direction, and after months of foot dragging, Intel seeks to stay discovery once again.

This case is not the case against AMD. A simple comparison of the two complaints show that Plaintiffs' allegations against Intel are substantially different. Indeed, it appears that among all major CPU manufacturers, only Intel CPUs suffer from the Meltdown and Foreshadow exploits because only Intel implemented speculative execution and memory access protection in a way to allow that prioritized speed over user's privileged information. *Id.* ¶¶ 265-266.

Indeed, Plaintiffs' Complaint does not suffer from the deficiencies the court in AMD highlighted, and Intel's invitation to the Court to revisit its earlier decision on the stay request based on different allegations against another manufacturer is vacuous. A comparison of the AMD court's decision dismissing the complaint with leave to amend to Plaintiffs' Complaint here shows that Plaintiffs have pled their claims and answered all the concerns/issued raised by the AMD court.

This Court will decide Intel's motion to dismiss based on the allegation in Plaintiffs' Complaint at the appropriate time. This Court denied Intel's previous request to stay discovery, and there is no reason to revisit that decision now.

For all of the foregoing reasons, Plaintiffs respectfully request that the Court direct Intel to commence production in accordance with the proposed TAR Protocol, set a near term deadline for exchange of Rule 26(a)(1) disclosures to be exchanged, and direct Intel to produce technical documents regarding the design, development, and engineering of Intel's implementations of speculative execution, memory access protection, and supporting technologies in its CPUs.

Respectfully submitted,

Christopher A. Seeger, Esq.

Rosemary M. Rivas, Esq.